Network Detection Theory and Performance

Steven T. Smith*, *Senior Member, IEEE*, Kenneth D. Senne*, *Life Fellow, IEEE*,
Scott Philips*, Edward K. Kao*†, and Garrett Bernstein*

arXiv:1303.5613v1 [cs.SI] 22 Mar 2013

***Abstract*—Network detection is an important capability in many areas of applied research in which data can be represented as a graph of entities and relationships. Oftentimes the object of interest is a relatively small subgraph in an enormous, potentially uninteresting background. This aspect characterizes network detection as a "big data" problem. Graph partitioning and network discovery have been major research areas over the last ten years, driven by interest in internet search, cyber security, social networks, and criminal or terrorist activities. The specific problem of network discovery is addressed as a special case of graph partitioning in which membership in a small subgraph of interest must be determined. Algebraic graph theory is used as the basis to analyze and compare different network detection methods. A new Bayesian network detection framework is introduced that partitions the graph based on prior information and direct observations. The new approach, called space-time threat propagation, is proved to maximize the probability of detection and is therefore optimum in the Neyman-Pearson sense. This optimality criterion is compared to spectral community detection approaches which divide the global graph into subsets or communities with optimal connectivity properties. We also explore a new generative stochastic model for covert networks and analyze using receiver operating characteristics the detection performance of both classes of optimal detection techniques.**

I. Introduction

Network detection is a special class of the more general graph partitioning (GP) problem in which the binary decision of membership or non-membership for each graph vertex must be determined. This detection problem and more generally GP are of fundamental and practical importance in graph theory and its applications (Figure 1). The detected subgraph comprises all vertices declared to be members. The very definition of membership will lead to specific network detection algorithms.

Graph partitioning is an NP-hard problem; however, semidefinite programming (SDP) relaxation applies to many cases, offering both practical and oftentimes theoretically attractive approximation to GP. [32], [56] In general, practical GP approaches exploit a variety of global and local connectivity properties to divide a graph into many subgraphs. Decreasing algorithmic complexity is achieved in certain domains that may be cast as quadratic optimization problems (yielding eigenvalue- or spectral-based methods), or simple sets of linear equations. One important network detection approach, called community detection, divides the global graph into subsets or communities based on optimizing a specific connectivity measure that is chosen depending upon the application. This paper presents a new Bayesian network detection approach called space-time threat propagation [40], [47] that is shown to optimize the probability of network detection in a Neyman-Pearson sense given prior information and/or direct observations, i.e. detection probability is maximized given a fixed false alarm a.k.a. false positive probability. This is an important property because it provides a practical optimum algorithm in many settings (satisfying a set of assumptions detailed later in the paper), and it provides a performance bound on detection performance. Remarkably, the two apparently different optimal network detection approaches are related to each other using insights from algebraic graph theory. Converse to other research on network detection, rather than using the network to detect signals, [3], [10], [27] the signal of interest in this paper *is* the signal to be detected. In this sense the paper is also related to work on so-called manifold learning methods, [5], [8], [12] although the network to be detected is a subgraph of an existing network, and therefore the methods described here belong to a class of network anomoly detection [8] as well as maximimum-likelihood methods for network detection. [17] Both spectral-based and Neyman-Pearson network detection methods are described and analyzed below in Sections III and IV. Furthermore, network detection performance is assessed using a new stochastic blockmodel [2] for small, dynamic foreground networks embedded within a

Manuscript received 2013.

*MIT Lincoln Laboratory, Lexington, MA 02420; { stsmith, senne, garrett.bernstein }@ll.mit.edu, edwardkao@fas.harvard.edu

†Department of Statistics, Harvard University; Cambridge MA USA 02138

*This work is sponsored by the Assistant Secretary of Defense for Research & Engineering under Air Force Contract FA8721-05-C-0002. Opinions, interpretations, conclusions and recommendations are those of the author and are not necessarily endorsed by the United States Government.

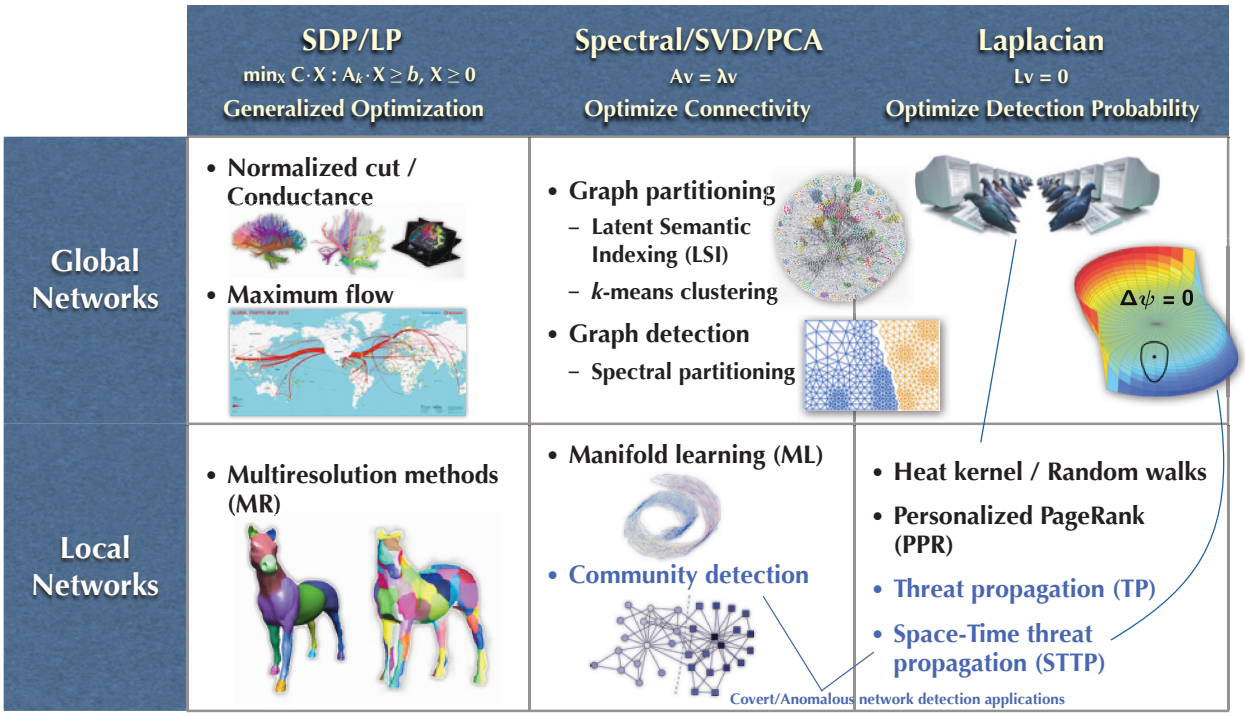

Fig. 1. Network detection algorithm taxonomy. This paper focuses on local spectral and harmonic methods for network detection. (Images used from various sources; clockwise from upper left [6, 48, 24, 38, 23, 44, 58]).

large background.

A. Covert Networks

Detection of network communities is most likely to be effective if the communities exhibit high levels of connection activity. However, the covert networks of interest to many applications are unlikely to cooperate with this optimistic assumption. Indeed, a "fully connected network ... is an unlikely description of the enemy insurgent order of battle." [50] A clandestine or covert community is more likely to appear cellular and distributed. [7] Communities of this type can be represented with "small world" models. [43] The covert networks of interest in this paper exist to accomplish nefarious, illegal, or terrorism goals, while "hiding in plain sight." [26], [57] Covert networks necessarily adopt operational procedures to remain hidden and robustly adapt to losses of parts of the network. For example, during the Algerian Revolution the FLN's Autonomous Zone of Algiers (Z.A.A.) military command was "carefully kept apart from other elements of the organization, the network was broken down into a number of quite distinct and compartmented branches, in communication only with the network chief," allowing Z.A.A. leader Yassef Saadi to command "within 200 yards from the office of the [French] army commandant ... and remain there several months[.]" [49] Krebs' reconstruction of the 9/11 terrorist network details the strategy for keeping cell members distant from each other and from other cells and notes bin Laden's description of this organization: "those ... who were trained to fly didn't know the others. One group of people did not know the other group." [30] A covert network does not have to be human to be nefarious; the widespread Flashback malware attack on Apple's OS X computers employed switched load balancing between servers to avoid detection, [14] mirroring the Z.A.A.'s "tree" structure for robust covert network organization.

In order to accomplish its goals the covert network must judiciously use "transitory shortcuts." [53] For example, in the 9/11 terrorism operation, after coordination meetings connected distant parts of the network, the "cross-ties went dormant." [30] It is during these occasional bursts of connection activity that a covert community may be most vulnerable to detection. [50]

Network detection is predicated on the existence of observations of network relationships. In this paper the focus will be on observations of network activities using Intelligence, Surveillance, and Reconnaissance (ISR) sensors, such as Wide-Area Motion Imagery (WAMI). Covert networks engaged in terrorist attacks with Improvised Explosive Devices (IEDs) comprise loosely connected cells with various functions, such as finance, planning, operations, logistics, security, and propaganda.

In this paper a new model of covert threat for detection analysis that accounts for the realities of dynamic foreground networks in large backgrounds is a specially adapted version of a mixed membership stochastic blockmodel. [2] The terrorist cells of interest are embedded into a background consisting of many "neutral" communities, that represent business, homes, industry, religion, sports, etc. Because in real life people wear different "hats" depending upon on the communities with which they interact, their proportions of membership in multiple communities (lifestyles) can be adjusted to control the occasional coordination between the foreground and background networks. The new generative blockmodel approach introduced in Section IV-A leads to a analytically tractable tool with sufficient parameters to exhibit realistic coordinated activity levels and interactions.

B. Observability and Detectability

The connections (edges) between nodes of a network are observable only when they are active. This implies that there are two basic strategies for detecting a covert threat: (1) subject-based Bayesian models that correlate a priori information or observations of the observed network connections; (2) pattern-based (predictive) methods that look for known patterns of organization/behavior to infer nefarious activity. [26], [41] Subject-based methods follow established principles of police investigations to accrue evidence based upon observed connections and historical data. The dependency of predictive methods on known patterns, however, makes them difficult to apply to rare and widely different covert threats: "there are no meaningful patterns that show what behavior indicates planning or preparation for terrorism." [26] The real-world consequences of applying an inappropriate model to detect a threat may include an unacceptable number of false positives and an erosion of individual privacy rights and civil liberties. [26], [41]

As described above, the subject of community detection in graphs has experienced extensive research during the last ten years. [21], [25], [29], [38], [39] Nevertheless, there are few closed-form results that quantify the limits of detectability of specific types of networks in representative backgrounds. Fully connected networks (cliques) have received special attention: there is a recent result which confirms in closed form using random matrix theory the previously observed phase transition of detectability for sufficiently small cliques [20], [31], [37] or dense subgraphs [4].

In this paper we use the proposed generative stochastic threat model with Monte Carlo detection performance analysis. The detection methodologies under investigation here include the spectral-based and Neyman-Pearson techniques discussed above in the Introduction.

II. Algebraic Graph Theory

A graph $G = (V, E)$ is defined by two sets, the vertices V of G, and the edges $E \subset [V]^2 \subset 2^V$ of G, in which $[V]^2$ denotes the set of 2-element subsets of V. [13] For example, the sets $V = \{\,1,\ 2,\ 3\,\}$, $E = \big\{\,\{1,\,2\},\ \{2,\,3\}\,\big\}$ describe a simple graph with undirected edges between vertices 1 and 2, and 2 and 3: $\textcircled{1}\!\!-\!\!\textcircled{2}\!\!-\!\!\textcircled{3}$. The *adjacency matrix* $\mathbf{A} = \mathbf{A}(G)$ of G is the $\{0,1\}$-matrix with $\mathbf{A}_{ij} = 1$ iff $\{\,i, j\,\} \in E$. In the example, $\mathbf{A} = \left(\begin{smallmatrix} 0 & 1 & 0 \\ 1 & 0 & 1 \\ 0 & 1 & 0 \end{smallmatrix}\right)$. Because simple graphs are undirected, their adjacency matrix is necessarily symmetric. The *degree matrix* $\mathbf{D} = \mathrm{Diag}(\mathbf{A}\cdot\mathbf{1})$ is the diagonal matrix of the vector of degrees of all vertices, where $\mathbf{1} = (1, \ldots, 1)^{\mathrm{T}}$ is the vector of all ones.

Many important applications involve an orientation between vertices, defined by an orientation map $\sigma\colon [V]^2 \to V \times V$ (the ordered Cartesian product of V with itself) in which the first and second coordinates are called the initial and terminal vertices, respectively. The corresponding *directed graph* is denoted G^σ or, by abuse of notation, simply G. The preceding example with orientation map $\sigma(\{1,\,2\}) = (2, 1)$, $\sigma(\{2,\,3\}) = (2, 3)$ yields the directed graph $\textcircled{1}\!\!\leftarrow\!\!\textcircled{2}\!\!\rightarrow\!\!\textcircled{3}$. The *incidence matrix* $\mathbf{B} = \mathbf{B}(G^\sigma)$ of the oriented graph G^σ is the $(0, \pm 1)$-matrix of size $\#V$-by-$\#E$ with $\mathbf{B}_{ie} = -1$ if i is an initial vertex of $\sigma(e)$, 1 if i is a terminal vertex of $\sigma(e)$, and 0 otherwise. In the example, $\mathbf{B} = \left(\begin{smallmatrix} 1 & 0 \\ -1 & -1 \\ 0 & 1 \end{smallmatrix}\right)$. In the study of homology in algebraic topology, the incidence matrix is recognized as the boundary operator on graph edges. It encodes differences between vertices and plays an important role in the analysis of network detection algorithms through the so-called graph Laplacian, which appears in three forms. The *unnormalized Laplacian matrix* or *Kirchhoff matrix* of a graph G is the matrix

$$\mathbf{Q} = \mathbf{Q}(G) = \mathbf{B}\mathbf{B}^{\mathrm{T}} = \mathbf{D} - \mathbf{A}, \tag{1}$$

where $\mathbf{B}(G^\sigma)$ is the incidence matrix of an oriented graph $\mathbf{G}^\sigma$ with (arbitrary) orientation σ, and $\mathbf{A}(G)$ and $\mathbf{D}(G)$ are, respectively, the adjacency and degree matrices of G. In the example, $\mathbf{Q} = \left(\begin{smallmatrix} 1 & -1 & 0 \\ -1 & 2 & -1 \\ 0 & -1 & 1 \end{smallmatrix}\right)$. The (normalized) *Laplacian matrix*

$$\mathbf{L} = \mathbf{D}^{-1/2}\mathbf{Q}\mathbf{D}^{-1/2} = \mathbf{I} - \mathbf{D}^{-1/2}\mathbf{A}\mathbf{D}^{-1/2} \tag{2}$$

is a matrix congruence of the Kirchhoff matrix $\mathbf{Q}$ scaled by the square-root of the degree matrix $\mathbf{D}^{1/2}$. The *generalized* or *asymmetric Laplacian matrix*

$$\mathbf{Ł} = \mathbf{D}^{-1/2}\mathbf{L}\mathbf{D}^{1/2} = \mathbf{D}^{-1}\mathbf{Q} = \mathbf{I} - \mathbf{D}^{-1}\mathbf{A} \tag{3}$$

is a similarity transformation of the Laplacian matrix. In the example, $\mathbf{L} = \begin{pmatrix} 1 & -2^{-1/2} & 0 \\ -2^{-1/2} & 1 & -2^{-1/2} \\ 0 & -2^{-1/2} & 1 \end{pmatrix}$ and $\mathbf{Ł} = \begin{pmatrix} 1 & -1 & 0 \\ -2^{-1} & 1 & -2^{-1} \\ 0 & -1 & 1 \end{pmatrix}$. The latter example is immediately recognized as a discretization of the second derivative $-d^2/dx^2$, i.e. the negative of the 1-d Laplacian operator $\Delta = \partial^2/\partial x^2 + \partial^2/\partial y^2 + \cdots$ that appears in numerous physical applications. (This sign is the convention used in graph theory.) The asymmetric Laplacian $\mathbf{Ł} = \mathbf{I} - \mathbf{D}^{-1}\mathbf{A}$ plays an important role in mean-value theorems involving solutions to Laplace's equation $\mathbf{Łv} = \mathbf{0}$, which will be seen to be the motivating equation behind several network detection algorithms.

The connection between the incidence and Laplacian matrices and physical applications is made through Green's first identity, which equates the continuous Laplacian operator Δ in terms of the vector gradient $\nabla = (\partial/\partial x, \partial/\partial y, \ldots)^{\mathrm{T}}$ and motivates the definition $\mathbf{Q} = \mathbf{B}\mathbf{B}^{\mathrm{T}}$ of the graph Laplacian. Given two arbitrary "test" functions $f(\mathbf{x})$ and $g(\mathbf{x})$ on a bounded domain $\Omega \subset \mathbb{R}^n$ with boundary $\partial\Omega$ and inner product $\langle , \rangle$, Green's first identity asserts,

$$\int_\Omega g\Delta f \, dV = -\int_\Omega \langle \nabla g, \nabla f \rangle \, dV + \int_{\partial\Omega} g\langle \nabla f, \mathbf{n} \rangle \, dS, \tag{4}$$

where dV and $\mathbf{n}\, dS$ are the volume and directed surface differentials—this formula generalizes immediately to Riemannian manifolds. Applying the finite element method to this continuous equation yields a graph arising from, say, Delaunay triangulation and a matrix equation involving the graph Laplacian matrix $\mathbf{L}$ [from $g\Delta f$ in Eq. (4)] and the normalized outer product $\mathbf{D}^{-1/2}\mathbf{B}\mathbf{B}^{\mathrm{T}}\mathbf{D}^{-1/2}$ of the incidence matrix [from $\langle \nabla g, \nabla f \rangle$ in Eq. (4)]. This illustrates that the graph Laplacian is the standard Laplacian of physics and mathematics, a connection that explains many theoretical and performance advantages of the normalized Laplacian over the Kirchhoff matrix across applications. [11], [45], [52], [54], [55]

The most important property of the Laplacian matrix is that the constant vector $\mathbf{1} = (1, \ldots, 1)^{\mathrm{T}}$ is in the kernel of the Laplacian,

$$\mathbf{Q1} = \mathbf{0}; \qquad \mathbf{Ł1} = \mathbf{0}, \tag{5}$$

i.e. $\mathbf{1}$ is an eigenvector of $\mathbf{Q}$ and $\mathbf{Ł}$ whose eigenvalue is zero. This property is the reason for the mean-value property of harmonic functions, as well as the fact that the only bounded harmonic functions on an unbounded domain are necessarily constant, which will play an important role in optimum network detection. This is a key fact because many network detection algorithms involve solutions to Laplace's equation, however this constant solution does not distinguish between vertices at all, a deficiency that may be resolved in a variety of ways, yielding a family of network detection algorithms. Furthermore, the geometric multiplicity of the zero eigenvalue equals the number of connected components of the graph, though because a connected graph is implicit for the subgraph detection problem, we may assume that the kernel of the graph Laplacian is simply the one-dimensional subspace $(\mathbf{1}) = \{ \alpha\mathbf{1} : \alpha \in \mathbb{R} \}$.

III. Optimum Network Detection

Two different optimality criteria are used for the two different strategies of network detection: various connectivity metrics are used for predictive methods, and detection performance is used for subject-based methods. Detection optimality means, as usual, optimality in the Neyman-Pearson sense in which the probability of detection is maximized at a fixed false alarm rate. In the context of networks, the probability of detection (PD) refers to the fraction of vertices detected belonging to the threat subgraph, and the probability of false alarm (PFA) refers to the fraction of non-threat vertices detected. As in classical detection theory, [51] the optimal detector is a threshold of the log-likelihood ratio (LLR), and a new Bayesian framework for network detection is developed in this section. The distinction between classical detection theory and network detection theory is not in the form of the optimal detector—the log-likelihood ratio—but in distinct mathematical formulations. Whereas linear algebra is the foundation for classical detection theory, algebraic graph theory [22] is the foundation for network detection. It follows that understanding the theory, algorithms, and results of network detection requires an introduction of some basic concepts from algebraic graph theory, especially the graph Laplacian and spectral analysis of graphs. Familiarization with these objects provides a common framework of comparing apparently unrelated network detection algorithms and provides deep insights into basic problems in network detection theory.

A. Spectral-Based Community Detection

Efficient graph partitioning algorithms and analysis appeared in the 1970s with Donath and Hoffman's eigenvalue-based bounds for graph partitioning [15] and Fiedler's connectivity analysis and graph partitioning algorithm [18], [19] which established the connection between a graph's algebraic properties and the spectrum of its Kirchhoff Laplacian matrix $\mathbf{Q} = \mathbf{D} - \mathbf{A}$ [Eq. (1)]. The spectral methods in this section solve the graph

partitioning problem by optimizing various subgraph connectivity properties.

The *cut size* of a subgraph—the number of edges necessary to remove to separate the subgraph from the graph—is quantified by the quadratic form $\mathbf{s}^{\mathrm{T}}\mathbf{Q}\mathbf{s}$, where $\mathbf{s} = (\pm 1, \ldots, \pm 1)^{\mathrm{T}}$ is a ± 1-vector who entries are determined by subgraph membership. [42] Minimizing this quadratic form over $\mathbf{s}$, whose solution is an eigenvalue problem for the graph Laplacian, provides a network detection algorithm based on the model of minimal cut size. However, there is a paradox in the application of spectral methods to network detection: the smallest eigenvalue of the graph Laplacian $\lambda_0(\mathbf{Q}) = 0$ corresponds to the eigenvector $\mathbf{1}$ constant over all vertices, which fails to discriminate between subgraphs. Intuitively this degenerate constant solution makes sense because the two subgraphs with minimal (zero) subgraph cut size are the entire graph itself ($\mathbf{s} \equiv \mathbf{1}$), or the null graph ($\mathbf{s} \equiv -\mathbf{1}$). This property manifests itself in many well-known results from complex analysis, such as the maximum principle.

Fiedler showed that if rather the eigenvector $\boldsymbol{\xi}_1$ corresponding to the second smallest eigenvalue $\lambda_1(\mathbf{Q})$ of $\mathbf{Q}$ is used (many authors write $\lambda_1 = 0$ and λ_2 rather than the zero offset indexing $\lambda_0 = 0$ and λ_1 used here), then for every nonpositive constant $c \leq 0$, the subgraph whose vertices are defined by the threshold $\boldsymbol{\xi}_1 \geq c$ is necessarily connected. This algorithm is called *spectral detection*. Given a graph G, the number $\lambda_1(\mathbf{Q})$ is called the *Fiedler value* of G, and the corresponding eigenvector $\boldsymbol{\xi}_1(\mathbf{Q})$ is called the *Fiedler vector*. Completely analogous with comparison theorems in Riemannian geometry that relate topological properties of manifolds to algebraic properties of the Laplacian, many graph topological properties are tied to its Laplacian. For example, the graph's diameter D and the minimum degree $d_{\min}$ provide lower and upper bounds for the Fiedler value $\lambda_1(\mathbf{Q})$: $4/(nD) \leq \lambda_1(\mathbf{Q}) \leq n/(n-1)\cdot d_{\min}$. [36] This inequality explains why the Fiedler value is also called the *algebraic connectivity*: the greater the Fiedler value, the smaller the graph diameter, implying greater graph connectivity. If the normalized Laplacian $\mathbf{L}$ of Eq. (2) is used, the corresponding inequality involving the generalized eigenvalue $\lambda_1(\mathbf{L}) = \lambda_1(\mathbf{Q}, \mathbf{D})$ involves the graph's diameter D and volume V: $1/(DV) \leq \lambda_1(\mathbf{L}) \leq n/(n-1)$. [11]

Because in practice spectral detection with its implicit assumption of minimizing the cut size oftentimes does not detect intuitively appealing subgraphs, Newman introduced the alternate criterion of subgraph "modularity" for subgraph detection. [38] Rather than minimize the cut size, Newman proposes to maximize the subgraph connectivity relative to background graph connectivity, which yields the quadratic maximization problem $\max_{\mathbf{s}} \mathbf{s}^{\mathrm{T}}\mathbf{M}\mathbf{s}$, where $\mathbf{M} = \mathbf{A} - V^{-1}\mathbf{d}\mathbf{d}^{\mathrm{T}}$ is Newman's *modularity matrix*, $\mathbf{A}$ is the adjacency matrix, $(\mathbf{d})_i = d_i$ is the degree vector, and $V = \mathbf{1}^{\mathrm{T}}\mathbf{d}$ is the graph volume. [38] Newman's modularity-based graph partitioning algorithm, also called community detection, involves thresholding the values of the principal eigenvector of $\mathbf{M}$. Miller et al. [33]–[35] also consider thresholding arbitrary eigenvectors of the modularity matrix, which by the Courant minimax principle biases the Newman community detection algorithm to smaller subgraphs, a desirable property for many applications. They also outline an approach for exploiting observations within the spectral framework. [33]

B. Neyman-Pearson Subgraph Detection

Network detection of a subgraph within a graph $G = (V, E)$ of order n is treated as n independent binary hypothesis tests to decide which of the graph's n vertices does not belong (null hypothesis H_0) or belongs (hypothesis H_1) to the network. Maximizing the probability of detection (PD) for a fixed probability of false alarm (PFA) yields the Neyman-Pearson test involving the log-likelihood ratio of the competing hypothesis. We will derive this test in the context of network detection, which both illustrates the assumptions that ensure detection optimality, as well as indicates practical methods for computing the log-likelihood ratio test and achieving an optimal network detection algorithm. It will be seen that a few basic assumptions yield an optimum test involving the graph Laplacian, which allows comparison of Neyman-Pearson testing to several other network detection methods whose algorithms are also related to the properties of the Laplacian.

Assume that each vertex $v \in V$ has an unknown $\{0, 1\}$-valued property Θ_v which is considered to be "threat" or "non-threat" at v, and that there exists an observation vector $\mathbf{z}\colon \{v_{i_1}, \ldots, v_{i_k}\} \subset V \to M \subset \mathbb{R}^k$ from k vertices to a measurement space M. For example, a direct observation of threat at vertex v may be represented by the observation $z(v) \equiv 1$. It is assumed that the observation $z(v)$ at v and the threat Θ_v at v are not independent, i.e. $f\big(z(v)|\Theta_v\big) \neq f\big(z(v)\big)$, so that there is positive mutual information between $z(v)$ and Θ_v. The probability density $f\big(z(v)|\Theta_v\big)$ is called the *observation model*, which in this paper is treated as a simple $\{0, 1\}$-valued model $\delta\big(z(v) - \Theta_v\big)$. Though the threat network hypotheses are being treated here independently at each vertex, this framework allows for more sophisticated global models that include hypotheses over two or more vertices.

An optimum hypothesis test is now derived for the presence of a network given a set of observations $\mathbf{z}$. Optimality is defined in the Neyman-Pearson sense in which the probability of detection is maximized at a constant false alarm rate (CFAR). As usual, [51] the derivation of the optimum test involves the procedure of Lagrange multipliers. For the general problem of network detection of a subgraph within graph G of order n, the decision of which of the 2^n hypothesis $\mathbf{\Theta} = (\Theta_{v_1}, \ldots, \Theta_{v_n})^{\mathrm{T}}$ to choose involves a 2^n-ary multiple hypothesis test over the measurement space of the observation vector $\mathbf{z}$, and an optimal test involves partitioning the measurement space into 2^n regions yielding a maximum PD. This NP-hard general combinatoric problem is clearly computationally and analytically intractable; however, the general 2^n-ary multiple hypothesis test may be greatly simplified by treating it as n independent binary hypothesis tests.

At each vertex $v \in G$ and unknown threat $\Theta\colon V \to \{0, 1\}$ across the graph , consider the binary hypothesis test for the unknown value Θ_v,

$$\begin{aligned} &\mathrm{H}_0(\mathrm{v}): \quad \Theta_v = 0 \qquad \text{(vertex belongs to background)} \\ &\mathrm{H}_1(\mathrm{v}): \quad \Theta_v = 1 \qquad \text{(vertex belongs to subgraph).} \end{aligned} \tag{6}$$

Given the observation vector $\mathbf{z}\colon \{v_{i_1}, \ldots, v_{i_k}\} \subset V \to M \subset \mathbb{R}^k$ with observation models $f\big(z(v_{i_j})|\Theta_{v_{i_j}}\big)$, $j = 1, \ldots, k$, the PD and PFA are given by the integrals

$$\mathrm{PD} = \int_R f(\mathbf{z}|\Theta_v = 1)\, d\mathbf{z}, \tag{7}$$

$$\mathrm{PFA} = \int_R f(\mathbf{z}|\Theta_v = 0)\, d\mathbf{z}, \tag{8}$$

where $R \subset M$ is the detection region in which observations are declared to yield the decision $\Theta_v = 1$, otherwise Θ_v is declared to equal 0. The optimum Neyman-Pearson test uses the detection region R that maximizes PD at a fixed CFAR value PFA_0. Posing this optimization problem over R with the method of Lagrange multipliers applied to the function

$$\begin{aligned} &F(R, \lambda) = \mathrm{PD}(R) - \lambda\big(\mathrm{PFA}(R) - \mathrm{PFA}_0\big), \\ &= \int_R f(\mathbf{z}|\Theta_v = 1)\, d\mathbf{z} - \lambda\Big[\int_R f(\mathbf{z}|\Theta_v = 0)\, d\mathbf{z} - \mathrm{PFA}_0\Big] \\ &= \int_R \big[f(\mathbf{z}|\Theta_v = 1) - \lambda f(\mathbf{z}|\Theta_v = 0)\big]\, d\mathbf{z} + \lambda\, \mathrm{PFA}_0 \end{aligned} \tag{9}$$

yields two conditions to maximize $F(R, \lambda)$ over R and λ:

(i) $\lambda > 0$,

(ii) $z \in R \;\Leftrightarrow\; f(\mathbf{z}|\Theta_v = 1) - \lambda f(\mathbf{z}|\Theta_v = 0) > 0$.

The second property yields the *likelihood ratio* (LR) test,

$$\frac{f(\mathbf{z}|\Theta_v = 1)}{f(\mathbf{z}|\Theta_v = 0)} \underset{\mathrm{H}_0(v)}{\overset{\mathrm{H}_1(v)}{\gtrless}} \lambda \tag{10}$$

that maximizes the probability of detection. As will be shown in the next section, the numerator $f(\mathbf{z}|\Theta_v = 1)$ of Eq. (10) is easily computed using standard Bayesian analysis, leading to a "threat propagation" algorithm for $f(\Theta_v|\mathbf{z})$ and a connection to the Laplacian $\mathbf{Ł}(G)$ described in Section II, and the denominator $f(\mathbf{z}|\Theta_v = 0)$ is determined by prior background information or simply the "principle of insufficient reason" [28] in which this term is a constant.

Because the probability of detecting threat is maximized at each vertex, the probability of detection for the entire subgraph is also maximized, yielding an optimum Neyman-Pearson test under the simplification of treating the 2^n-ary multiple hypothesis testing problem as a sequence of n binary hypothesis tests. Summarizing, the probability of network detection given an observation $\mathbf{z}$ is maximized by computing $f(\Theta_v|\mathbf{z})$ using a Bayesian "threat propagation" method and applying a simple likelihood ratio test. The connectivity of the subgraph whose vertices exceed the threshold is assured by the maximum principle. Algorithms for computing $f(\Theta_v|\mathbf{z})$ are described next.

C. Space-Time Threat Propagation

Many important network detection applications, especially networks based on vehicle tracks and computer communication networks, involve directed graphs in which the edges have departure and arrival times associated with their initial and terminal vertices. Space-Time threat propagation is used compute the time-varying threat across a graph given one or more observations at specific vertices and times. [40], [47] In such scenarios, the time-stamped graph $G = (V, E)$ may be viewed as a *space-time graph* $G_T = (V \times T, E_T)$ where T is the set of sample times and $E_T \subset [V \times T]^2$ is an edge set determined by the temporal correlations between vertices at specific times. This edge set is application-dependent, but must satisfy the two constraints, (1) if $\big(u(t_k), v(t_l)\big) \in E_T$ then $(u, v) \in E$, and (2) temporal subgraphs $\big((u, v), E_T(u, v)\big)$ between any two vertices u and v are defined by a temporal model $E_T(u, v) \subset [T \coprod T]^2$. A concrete example for a specific dynamic model of threat propagation is provided below.

1) Temporal Threat Propagation: Given an observed threat at a particular vertex and time, we wish to compute the inferred threat across all vertices and all times. This computation is a straightforward application of Bayesian

analysis that results in the optimum Neyman-Pearson network detection test developed above as well as an efficient algorithm for computing this test. Given a vertex v, denote the threat at v and at time $t \in \mathbb{R}$ by the $\{0, 1\}$-valued stochastic process $\Theta_v(t)$, with value zero indicating no threat, and value unity indicating a threat. Denote the *probability of threat* at v at t by

$$\vartheta_v(t) \stackrel{\text{def}}{=} P\big(\Theta_v(t) = 1\big) = P\big(\Theta_v(t)\big). \tag{11}$$

The threat state at v is modeled by a finite-state continuous time Markov jump process between from state 1 to state 0 with Poisson rate λ_v. With this simple model the threat stochastic process $\Theta_v(t)$ satisfies the Itô stochastic differential equation,

$$d\Theta_v = -\Theta_v \, dN_v; \quad \Theta_v(0) = \theta_1, \tag{12}$$

where $N_v(t)$ is a Poisson process with rate λ_v defined for positive time, and simple time-reversal provides the model for negative times. Given an observed threat $z = \Theta_v(0) = 1$ at v at $t = 0$ so that $\vartheta_V(0) = 1$, the probability of threat at v under the Poisson process model (including time-reversal) is

$$\vartheta_v(t) = P\big(\Theta_v(t) | z = \Theta_v(0) = 1\big) = e^{-\lambda_v |t|}, \tag{13}$$

This stochastic model provides a Bayesian framework for inferring, or propagating, threat at a vertex over time given threat at a specific time. The function

$$K_v(t) = e^{-\lambda_v |t|} \tag{14}$$

of Eq. (13) is called the *space-time threat kernel* and when combined with spatial propagation provides a temporal model E_T for a space-time graph. A Bayesian model for propagating threat from vertex to vertex will provide a full space-time threat propagation model and allow for the application of the optimum maximum likelihood test of Eq. (10).

2) Spatial Threat Propagation: Propagation of threat from vertex to vertex is determined by tracks or connections between vertices. A straightforward Bayesian analysis yields nonlinear equations that determine the probability of threat at each vertex, and along with the assumptions of asymptotic independence and small probabilities these equations may be linearized and thereby easily analyzed and solved in regimes relevant to our applications.

The threat at vertex v at which a single track τ from vertex u arrives and/or departs at times t_τ^v and t_τ^u is determined by Eq. (13) and the (independent) event $v \leftarrow u$ that threat traveled along this track: $P\big(\Theta_v(t)\big) = \vartheta_v(t) = \vartheta_u(t_\tau^u) K_v(t - t_\tau^v) P(v \leftarrow u)$. There is a linear transformation

$$\begin{aligned} \vartheta_v(t) &= P(v \leftarrow u) K(t - t_\tau^v) \vartheta_u(t_\tau^u) \\ &= \int_{-\infty}^{\infty} P(v \leftarrow u) K(t - t_\tau^v) \delta(\sigma - t_\tau^u) \vartheta_u(\sigma) \, d\sigma \end{aligned} \tag{15}$$

from the threat probability at u to v. Discretizing time, the temporal matrix $\mathbf{K}_\tau^{uv}$ for the discretized operator has the sparse form

$$\mathbf{K}_\tau^{uv} = \Big(\mathbf{0} \ \ldots \ \mathbf{0} \, K(t_k - t_\tau^v) \, \mathbf{0} \ \ldots \ \mathbf{0} \Big), \tag{16}$$

where $\mathbf{0}$ represents an all-zero column, t_k represents a vector of discretized time, and the discretized function $K(t_k - t_\tau^v)$ appears in the column corresponding to the discretized time at t_τ^u. Threat propagating from vertex v to u along the same track τ is given by the comparable expression $\vartheta_u(t) = \vartheta_v(t_\tau^v) K(t - t_\tau^u)$, whose discretized linear operator $\mathbf{K}_\tau^{vu}$ takes the form

$$\mathbf{K}_\tau^{vu} = \Big(\mathbf{0} \ \ldots \ \mathbf{0} \, K(t_k - t_\tau^u) \, \mathbf{0} \ \ldots \ \mathbf{0} \Big) \tag{17}$$

[cf. Eq. (16)] where the nonzero column corresponds to t_τ^v. The sparsity of $\mathbf{K}_\tau^{uv}$ and $\mathbf{K}_\tau^{vu}$ will be essential for practical space-time threat propagation algorithms.

It will now be shown how threats arriving on other tracks from other vertices may be sequentially linearized. If the threat on the track τ from vertex u must be combined with an existing threat $\vartheta_v(t)$ at v, then the combined threat $\vartheta_v(t^\pm) = P\big(\Theta_v(t^\pm)\big)$ at v at time $t^\pm$ immediately after/before the track from u arrives/departs at time t is determined by the addition law of probability,

$$\begin{aligned} P\big(\Theta_v(t) \cup \Theta_u(t)(v \leftarrow u)\big) \\ = P\big(\Theta_v(t)\big) + P\big(\Theta_u(t)(v \leftarrow u)\big) \\ - P\big(\Theta_v(t) \cdot \Theta_u(t)(v \leftarrow u)\big). \end{aligned} \tag{18}$$

Under the two assumptions that (1) the threat events Θ_u and Θ_v at u and v are independent, asymptotically valid for large time differences relative to the Poisson time $\lambda_{v^*}^{-1}$ for an observation at vertex v^*, [47] and (2) the threat probabilities $P\big(\Theta_u(t)\big)$ and $P\big(\Theta_v(t)\big)$ are numerically small, Eq. (18) yields the linear approximation

$$\begin{aligned} \vartheta_v(t^\pm) &\approx P\big(\Theta_v(t)\big) + P\big(\Theta_u(t)(v \leftarrow u)\big) \\ &= \vartheta_v(t) + \vartheta_u(t) P(v \leftarrow u). \end{aligned} \tag{19}$$

Extending this analysis to multiple tracks and assuming that $P(v \leftarrow u)^{-1} \propto w(v)$ for some weight function $w \colon V \to \mathbb{R}$ of the vertices, e.g., the degree of each vertex, yields the *threat propagation equation*

$$\boldsymbol{\vartheta} = \mathbf{D}^{-1} \mathbf{A} \boldsymbol{\vartheta}, \tag{20}$$

where $\boldsymbol{\vartheta}$ is the (discretized) space-time vector of threat probabilities, the weighted space-time adjacency matrix

$$\mathbf{A}_{uv} = \begin{pmatrix} \mathbf{0} & \sum_l \mathbf{K}_{\tau_l}^{vu} \\ \sum_l \mathbf{K}_{\tau_l}^{uv} & \mathbf{0} \end{pmatrix} \tag{21}$$

is defined by Eq. (16), and $\mathbf{D}^{-1} = \operatorname{diag}\big(w(v_1)\mathbf{I}, \ldots, w(v_n)\mathbf{I}\big)$. Eq. (20), written as $\boldsymbol{\text{Ł}}\boldsymbol{\vartheta} = \mathbf{0}$, connects the asymmetric Laplacian matrix of Eq. (3) with threat propagation, the solution of which itself may be viewed as a boundary value problem with the harmonic operator $\boldsymbol{\text{Ł}}$.

Given a cue at vertices $v_{b_1}, \ldots, v_{b_C}$, the *harmonic space-time threat propagation equation* is

$$\begin{pmatrix} \boldsymbol{\text{Ł}}_{\mathrm{ii}} & \boldsymbol{\text{Ł}}_{\mathrm{ib}} \end{pmatrix} \begin{pmatrix} \boldsymbol{\vartheta}_{\mathrm{i}} \\ \boldsymbol{\vartheta}_{\mathrm{b}} \end{pmatrix} = \mathbf{0} \tag{22}$$

where the space-time Laplacian $\boldsymbol{\text{Ł}} = \begin{pmatrix} \boldsymbol{\text{Ł}}_{\mathrm{ii}} & \boldsymbol{\text{Ł}}_{\mathrm{ib}} \\ \boldsymbol{\text{Ł}}_{\mathrm{bi}} & \boldsymbol{\text{Ł}}_{\mathrm{bb}} \end{pmatrix}$ and the space-time threat vector $\boldsymbol{\vartheta} = \begin{pmatrix} \boldsymbol{\vartheta}_{\mathrm{i}} \\ \boldsymbol{\vartheta}_{\mathrm{b}} \end{pmatrix}$ have been permuted so that cued vertices are in the 'b' blocks (the "boundary"), non-cued vertices are in 'i' blocks (the "interior"), and the cued space-time vector $\boldsymbol{\vartheta}_{\mathrm{b}}$ is given. The *harmonic threat* is the solution to Eq. (22),

$$\boldsymbol{\vartheta}_{\mathrm{i}} = -\boldsymbol{\text{Ł}}_{\mathrm{ii}}^{-1}(\boldsymbol{\text{Ł}}_{\mathrm{ib}}\boldsymbol{\vartheta}_{\mathrm{b}}). \tag{23}$$

The space-time Laplacian of Eq. (3) is a directed Laplacian matrix, and that Eq. (22) is directly analogous to Laplace's equation $\Delta\varphi = 0$ given a fixed boundary condition. As discussed in the next subsection, the connection between space-time threat propagation and harmonic graph analysis also provides a link to spectral-based methods for network detection. The nonnegativity of the harmonic threat of Eq. (23) is guaranteed because the space-time adjacency matrix $\mathbf{A}$ and cued threat vector $\boldsymbol{\vartheta}_{\mathrm{b}}$ are both nonnegative. This highly sparse linear system may be solved by the biconjugate gradient method, which provides a practical computational approach that scales well to graphs with thousands of vertices and thousands of time samples, resulting in space-time graphs of order ten million or more. In practice, significantly smaller subgraphs are encountered in applications such as threat network discovery [46], for which linear solvers with sparse systems are extremely fast.

Finally, a simple application of Bayes' theorem to the harmonic threat $\vartheta_v = f(\Theta_v|\mathbf{z})$ provides the optimum Neyman-Pearson detector [Eq. (10)] developed in Section III-B because

$$\begin{aligned} \frac{f(\mathbf{z}|\Theta_v = 1)}{f(\mathbf{z}|\Theta_v = 0)} &= \frac{f(\Theta_v = 1|\mathbf{z})}{f(\Theta_v = 0|\mathbf{z})} \cdot \frac{f(\Theta_v = 0)}{f(\Theta_v = 1)} \\ &= \frac{\vartheta_v}{f(\Theta_v = 0|\mathbf{z})} \cdot \frac{f(\Theta_v = 1)}{f(\Theta_v = 0)} \underset{\mathrm{H}_0(v)}{\overset{\mathrm{H}_1(v)}{\gtrless}} \lambda, \end{aligned} \tag{24}$$

results in a threshold of the harmonic space-time threat propagation vector

$$\boldsymbol{\vartheta} \underset{\mathrm{H}_0}{\overset{\mathrm{H}_1}{\gtrless}} \text{threshold}, \tag{25}$$

possibly weighted by a nonuniform null distribution $f(\Theta_v = 0|\mathbf{z})$, with the normalizing constant $f(\Theta_v = 1)/f(\Theta_v = 0)$ being absorbed into the detection threshold. This establishes, under the assumptions and approximations enumerated above, the detection optimality of harmonic space-time threat propagation.

D. Insights from Spectral Graph Theory

Each network detection algorithm above can be compared to each other by different approaches taken to address the problem posed by the (physical) fact that the smallest eigenvalue of the graph Laplacian is zero: $\mathbf{Q}\mathbf{1} = 0 \cdot \mathbf{1}$. Fiedler's spectral detection, which minimizes the network cut size, thresholds the eigenvector corresponding to the second smallest eigenvalue of the Laplacian—the Fiedler value. In contrast, community detection, which maximizes the subgraph connectivity relative to the background, recasts the objective of spectral detection resulting in a threshold of the principal or other eigenvectors of Newman's modularity matrix $\mathbf{M} = \mathbf{A} - V^{-1}\mathbf{d}\mathbf{d}^{\mathrm{T}}$. Alternatively, threat propagation, which maximizes the Bayesian probability of detection by computing the harmonic solution to Laplace's equation, $\boldsymbol{\text{Ł}}\boldsymbol{\vartheta} = \mathbf{0}$, but treats this as a boundary value problem with observations representing the boundary values and unknown values representing the interior.

E. Computational Complexity

Depending upon sparsity, the computational complexity of spectral methods ranges from $O(n \log n)$–$O(n^2)$ for principal eigenvector methods [38] to $O(n^2 \log n)$–$O(n^3)$ for methods that rely on full eigensolvers [33]–[35], with the lower cost exhibited with graphs whose average degree is over $\log n$, below which a random Erdős-Rényi graph is almost surely disconnected. [16] The cost of harmonic methods is about $O(n \log n)$–$O(n^2)$ for sparse matrix inversion and also depends upon the graph's sparsity. In practice, Arnoldi iteration can be used for sparse eigenvalue computation and the biconjugate gradient method can be used for sparse matrix inversion.

IV. Network Detection Performance

There are two ways to demonstrate network detection performance: empirical and theoretical, both of which depend on detailed knowledge of network behavior and

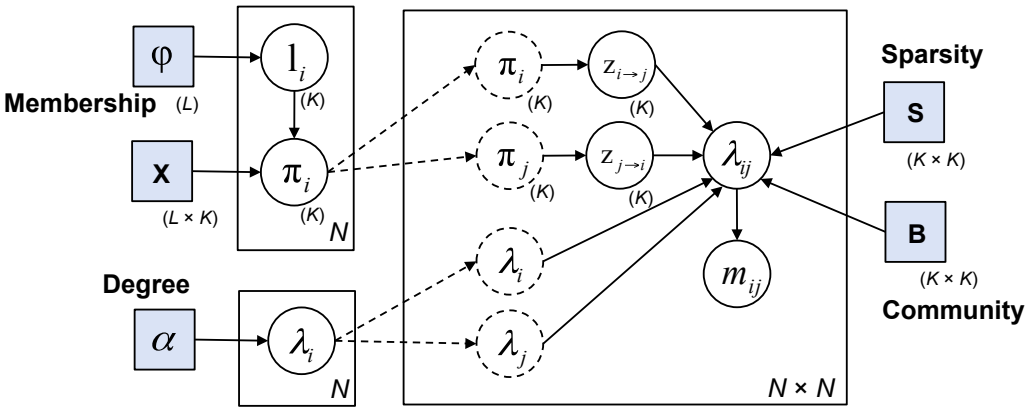

Fig. 2. Bayesian generative model for the network simulation with N nodes, K communities, and L "lifestyles" (distributions of community participation). Shaded squares are model parameters for tuning and circles are variables drawn during simulation.

dynamics. But full knowledge of real-world covert network behavior including relationships to the background network is, by design, extraordinarily rare or nonexistent, though partial information about many covert networks has been integrated over time [57]. Predicting performance of network detection methods requires details of the interconnectivity of both the foreground and background networks. Empirical detection performance is demonstrated using either a real-world or simulated dataset for which the truth is at least partially known, and theoretical performance predictions are derived based upon statistical assumptions about the foreground and background networks. To date, closed-form analytic performance predictions have been accomplished for very simple network models, i.e. cliques [20], [31], [37] or dense subgraphs [4] embedded within Erdős-Rényi backgrounds, and there are no theoretical results at all for space-time graphs or realistic models appropriate for covert networks. Therefore, realistic models are essential for performance analysis of network detection algorithms. There are two basic approaches to modeling networks: stochastic models, which attempt to capture the aggregate statistical properties of networks, and agent-based models, which attempt to describe specific behaviors. In general, stochastic models have greater tractability because they do not rely on the detailed description of actions or objectives of a specific network.

The empirical detection performance of the covert network detection algorithms described above will be computed using a Monte-Carlo analysis based upon a new stochastic blockmodel. Empirical performance predictions may be also based on a single dataset, oftentimes a practical necessity for real-world measurements. Detection performance for specific, real-world single datasets is illustrated in an accompanying paper. [59]

A. Covert Network Stochastic Blockmodel

To adhere with observed phenomenology of real-world networks, realistic network models should exhibit properties including connectedness, a power-law degree distribution (the "small world" property), membership-based community structure, sparsity, and temporal coordination. No one simple network model captures all these traits, e.g. Erdős-Rényi graphs can be almost surely connected, though do not exhibit a power-law density, power-law models such as R-MAT [9] do not exhibit a membership-based network structure, and mixed-membership stochastic blockmodels [2] do not include temporal coordination. To achieve a realistic network model possessing this range of properties, we propose a new statistical-based model with parameterized control over the generation of interactions between network nodes. The proposed model is depicted in Fig. 2 using plate notation.

1) Spatial Stochastic Blockmodel: The proposed model may be viewed as an aggregation of the several simpler models of which it is comprised: Erdős-Rényi (dominant at low degrees), [16] Chung-Lu (dominant at high degrees), [1] and a mixed-membership blockmodel that models community interactions. [2] The overall network model is approximated by each of the simpler models in the regime where the simple model dominates. The Erdős-Rényi model defines the overall sparsity and connectivity. The Chung-Lu model creates a power-law degree distribution empirically consistent with a broad range of real-world networks. The stochastic blockmodel creates distinct communities each with their own parameterized interaction models.

The space-time graph of the proposed mixed-membership stochastic blockmodel is determined by a connectivity model and temporal model. Let N be the total number of nodes, and K be the number of com-

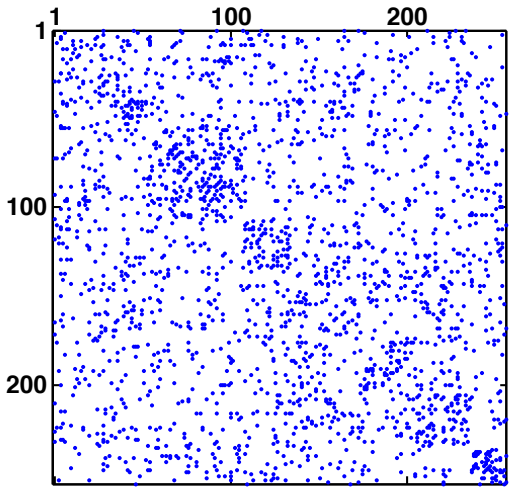

Fig. 3. Adjacency matrix of a stochastic blockmodel with a foreground subgraph whose intra-activity is 50% more than of all other subgraphs.

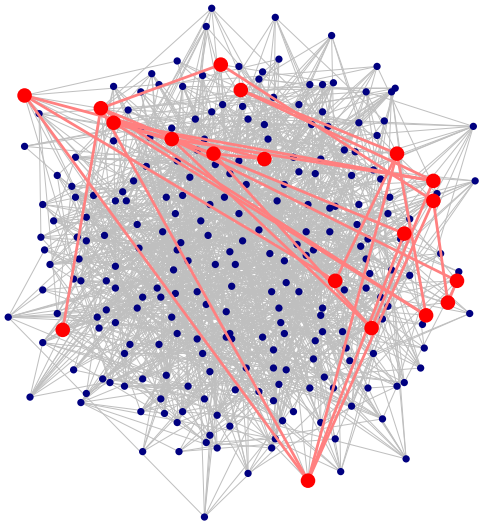

Fig. 4. Graph of the adjacency matrix shown in Figure 3. The foreground graph and intra subgraph edges are shown in red.

munities. Each node divides its time among at least one of the several K communities, and the number of ways in which a node distributes its time among the different communities is discretized into L distinct "lifestyles." Each node is assigned to a specific lifestyle. For example, nodes 1 and 3 may spend all their time in community 1, thereby sharing the same lifestyle, whereas node 2 may spend half its time in community 1 and half in community 2, and therefore occupies another lifestyle, and so forth. The rate λ_{ij} of interactions between nodes $i = 1$ and j is given by the product

$$\lambda_{ij} = I_{ij}^{\mathbf{S}} \cdot \frac{\lambda_i \lambda_j}{\sum_k \lambda_k} \cdot \mathbf{z}_{i \to j}^{\mathrm{T}} \mathbf{B} \mathbf{z}_{i \to j}, \tag{26}$$

where the first term $I_{ij}^{\mathbf{S}}$ represents the (modified) Erdős-Rényi model, the second term $\lambda_i \lambda_j / \left(\sum_k \lambda_k\right)$ represents the Chung-Lu model, and the third term $\mathbf{z}_{i \to j}^{\mathrm{T}} \mathbf{B} \mathbf{z}_{i \to j}$ represents the stochastic blockmodel.

At each node-to-node interaction, a random draw from a multinomial distribution determines the community to which each node belongs. The indicator function $I_{ij}^{\mathbf{S}}$ is a sparse K-by-K $(0,1)$-matrix whose entries are binomial random variables with probability $(\mathbf{S})_{ab}$ for node i in community a and node j in community b. An Erdős-Rényi sparsity model has $(\mathbf{S})_{ab} \equiv p$ for all communities a and b, whereas this modified sparsity model allows the possible of differing interaction rates within and across communities. The Chung-Lu term $\lambda_i \lambda_j / \left(\sum_k \lambda_k\right)$ is determined by the per-node expected degrees λ_i, $i = 1, \ldots, N$, which are themselves drawn from a power-law distribution of parameter $\boldsymbol{\alpha} \in \mathbb{R}^N$. The blockmodel term $\mathbf{z}_{i \to j}^{\mathrm{T}} \mathbf{B} \mathbf{z}_{i \to j}$ is determined by $\mathbf{B}$, a K-by-K matrix of the rate of interaction between communities, and $\mathbf{z}_{i \to j} \in \mathbb{R}^K$ is an indicator (01)-vector is the community to which node i belongs when interacting with node j. This community is the same over the entire simulation and is drawn from a multinomial over $\boldsymbol{\pi}_i \in \mathbb{R}^K$, node i's distribution over communities. Finally, the distribution of $\boldsymbol{\pi}$ is drawn from a Dirichlet r.v. with concentration parameter $\mathbf{l}_i^{\mathrm{T}} \mathbf{X}$. Node i's lifestyle, $\mathbf{l}_i$ is a multinomial draw with the lifestyle probability $\boldsymbol{\phi} \in \mathbb{R}^L$. The adjacency matrix and graph of this model are illustrated in Figs. 3 and 4 using an example with a mixed community with a higher level of activity for the foreground network.

2) Temporal Stochastic Blockmodel: The meeting times for each interaction are chosen independently of the spatial model. Real-world interactions are often coordinated, with many individuals arriving or leaving from a location at a set of pre-defined times. This behavior is parameterized by an average number of meeting times $\boldsymbol{\Psi} \in \mathbb{R}^K$ for each community. The simulated number of meeting times is a Poisson r.v. (offset by 1) with Poisson parameter $\boldsymbol{\Psi} - 1$. E.g. an expected number of meeting times $(\boldsymbol{\Psi})_k = 1$ (for community k) yields a constant Poisson r.v. of 1 meeting time (in Matlab, `poissrnd(0) = 0`), thereby yielding a community whose activities are tightly coordinated because there is only a single time for the members to meet. An expected number of meeting times $(\boldsymbol{\Psi})_k = 20$ yields a community whose activities are loosely coordinated because meetings may occur at any one of a number of times. The meetings times themselves are chosen uniformly over time, and each node arrives at the meeting time perturbed by a zero-mean Gaussian r.v. with a parameterized variance.

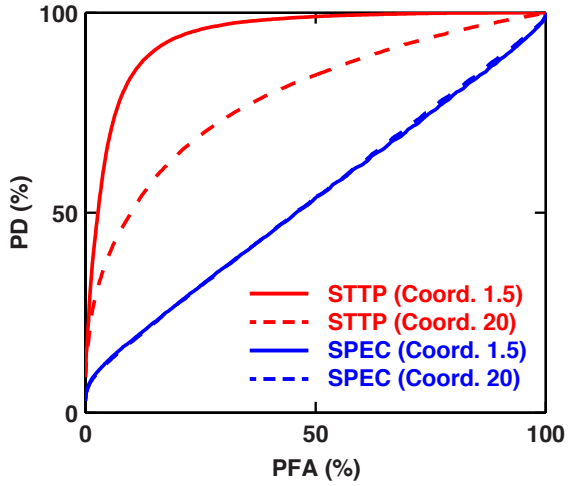

Fig. 5. Receiver operating characteristics versus forground coordination ($\mathbf{\Psi}_{\text{foreground}} = 1.5$, high coordination, and 20, low coordination) for space-time threat propagation (STTP) and spectral-based community detection (SPEC). The community activity level $\mathbf{S}_{\text{k}} = 1 \cdot \log N_{\text{k}}/N_{\text{k}}$ for all communities. [1000 Monte Carlo trials.]

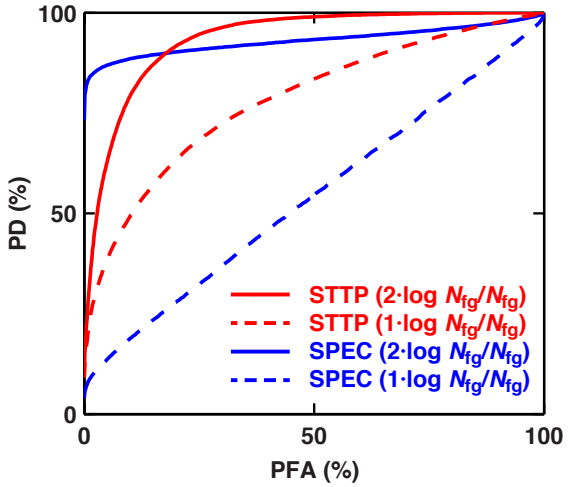

Fig. 6. Receiver operating characteristics versus foreground activity ($\mathbf{S}_{\text{fg}} = 2 \cdot \log N_{\text{fg}}/N_{\text{fg}}$, high activity, and $1 \cdot \log N_{\text{fg}}/N_{\text{fg}}$, baseline activity) for space-time threat propagation (STTP) and spectral-based community detection (SPEC). The foreground coordination level is specified by $\mathbf{\Psi}_{\text{fg}} = 20$ average number of meeting times. [1000 Monte Carlo trials.]

B. Network Detection Results

The detection performance of the network detection algorithms described above is presented in this section using empirical Monte Carlo results applied to the mixed-membership stochastic blockmodel. A space-time graph is chosen independently for each Monte Carlo trial. A set of baseline parameters is chosen to achieve realistic foreground and background networks of specific sizes, and excursions are performed on the parameters controlling foreground coordination and foreground activity. The performance metric is the standard receiver operating characteristic (ROC), which in the case of network detection is the probability of detection (measured as the percentage of true foreground nodes detected) versus the number or percentage of false alarms (the number of background nodes detected) as the detection threshold is varied. Perfect ROC performance is a 100% detection rate with a 0% false alarm rate, and the worst possible performance is a detection rate equal to chance, i.e. equal to the false alarm rate.

1) Baseline Model: A baseline model is used comprised of eleven lifestyles spanning ten communities. Two of the lifestyles are designated as foreground lifestyles and all others are "background." As detail above, each lifestyle has a propensity toward a different mix of community activity. The background lifestyles have a power-law distribution of membership over the background communities, which may be imagined to represent business, homes, industry, religion, sports, or other social interactions. Two distinct foreground lifestyles are used to model the compartmentalization of real-world covert networks. One foreground lifestyle associates uniformly across background communities, whereas the other foreground lifestyle has a strong association with a only small subset of background communities. These foreground lifestyles may be imagined to represent specialized functions or activities within the covert network. As in real life, the foreground lifestyles comprise only a tiny fraction of the entire population.

Interactions in which two nodes belong to the same community occur at a higher rate than interactions of nodes belonging to different communities. This is modeled by specifying that the block matrix $\mathbf{B}$ be diagonally dominant, perhaps strongly. Furthermore, real-world communities are not disconnected, thus for a community size of N_k, the diagonals of the Erdős-Rényi sparsity parameter matrix $\mathbf{S}_k$ must be at least $\log N_k/N_k$ to ensure that each community is almost surelyconnected [16]. Finally, covert networks necessarily have sparse—not clique-like—structure, thus the Erdős-Rényi sparsity parameter $\mathbf{S}$ for the covert network must also be low.

2) Detection versus Foreground Coordination and Activity: Two nominal values are chosen for foreground coordination and foreground activity, then both space-time threat propagation and spectral-based community detection algorithms are applied using a randomized cue over 1000 Monte Carlo trials. Fig. 5 shows the detection performance of both algorithms as the foreground coordination changes from a high of $\mathbf{\Psi}_{\text{fg}} = 1.5$ average number of meeting times to a low of $\mathbf{\Psi}_{\text{fg}} = 20$. As

predicted, the detection performance of space-time threat propagation improves as the temporal coordination of the foreground network increases. The optimality of this Bayesian network detector is predicated on temporal coordination, and decreased coordination makes the foreground network more difficult to detect. This example uses a constant baseline level of community activity (sparsity matrix $\mathbf{S}_{\rm fg} = 1 \cdot \log N_{\rm fg}/N_{\rm fg}$), thus the optimality assumption of high foreground activity made by spectral-based community detection algorithm is violated, and as expected this spectral algorithm does no better than chance for either coordination level. Fig. 6 shows the detection performance of both algorithms as the foreground activity changes from $\mathbf{S}_{\rm fg} = 1 \cdot \log N_{\rm fg}/N_{\rm fg}$ (baseline activity) to $\mathbf{S}_{\rm fg} = 2 \cdot \log N_{\rm fg}/N_{\rm fg}$ (high activity). The foreground coordination level is low, at $\mathbf{\Psi}_{\rm fg} = 20$, providing an example for which none of the basic algorithmic assumptions hold for either space-time threat propagation or spectral-based community detection. The low foreground activity results, $\mathbf{S}_{\rm fg} = 1 \cdot \log N_{\rm fg}/N_{\rm fg}$, are replicated in this figure from Fig. 5, in which STTP yields moderate detection performance and spectral-based community detection is no better than chance. At high foreground activity the foreground network is detectable at by both spectral-based community detection and space-time threat propagation.

V. Conclusions

The problem of covert network detection is analyzed from the perspectives of graph partitioning and algebraic graph theory. Network detection is addressed as a special case of graph partitioning in which membership in a small subgraph of interest must be determined, and a common framework is developed to analyze and compare different network detection methods. A new Bayesian network detection framework called space-time threat propagation is introduced that partitions the graph based on prior information and direct observations. Space-time threat propagation is shown to be optimum in the Neyman-Pearson sense subject to the assumption that threat networks are connected by edges temporally correlated to a cue or observation. Bayesian space-time threat propagation is interpreted as the solution to a harmonic boundary value problem on the graph, in which a linear approximation to Bayes' rule determines determines the unknown probability of threat on the uncued nodes (the "interior") based on threat observations at cue nodes (the "boundary"). This new method is compared to well-known spectral methods by examining competing notions of network detection optimality. Finally, a new generative mixed-membership stochastic blockmodel is introduced for performance prediction network detection algorithms. The parameterized model combines key real-world aspects of several random graph models: Erdős-Rényi for sparsity and connectivity, Chung-Lu for power-law degree distributions, and a mixed-membership stochastic blockmodel for distinctive community-based interaction and dynamics. This model is used to compute empirical detection performance results for the detection algorithms described in the paper as both foreground coordination and activity levels are varied. Though the results in the paper are empirical, it is our hope that both the paper's analytic results and performance modeling will be useful in future closed-form analysis of real-world covert network detection problems.

References

[1] W. Aiello, F. Chung, and L. Lu. "A random graph model for power law graphs," *Experimental Mathematics* **10** (1) : 53–66 (2001).

[2] E. M. Airoldi, D. M. Blei, S. E. Fienberg, and E. P. Xing. "Mixed-membership stochastic blockmodels," *JMLR* **9** : 1981–2014 (2008).

[3] M. Alanyali, S. Venkatesh, O. Savas, and S. Aeron. "Distributed Bayesian hypothesis testing in sensor networks," in *Proc. 2005 American Control Conf.* Boston MA, pp. 5369–5374 (2004).

[4] E. Arias-Castro and N. Verzelen. "Community Detection in Random Networks," `arXiv:1302.7099 [math.ST]`. Accessed 18 March 2013. ⟨http://arxiv.org/abs/1302.7099⟩.

[5] M. Belkin and P. Niyogi. "Laplacian eigenmaps for dimensionality reduction and data representation," *Neural Computation* **15** : 1373–1396 (2003).

[6] A. Brun, H. Knutsson, H. J. Park, M. E. Shenton, and C.-F. Westin. "Clustering fiber tracts using normalized cuts," in *Proc. Medical Image Computing and Computer-Assisted Intervention (MICCAI 04)* (2004). Accessed 3 September 2012. ⟨http://lmi.bwh.harvard.edu/papers/papers/brunMICCAI04.html⟩.

[7] K. Carley. "Estimating vulnerabilities in large covert networks," in *Proc. 16th Intl. Symp. Command and Control Research and Tech. (ICCRTS)*. (San Diego, CA) (2004).

[8] K. M. Carter, R. Raich, and A. O. Hero III. "On Local Intrinsic Dimension Estimation and Its Applications," *IEEE Trans. Signal Processing* **58** (2) : 650–663 (2010).

[9] D. Chakrabarti, Y. Zhan, and C. Faloutsos. "R-MAT: A Recursive Model for Graph Mining," in *Proc. 2004 SIAM Intl. Conf. Data Mining*. (2004).

[10] J.-F. Chamberland and V. V. Veeravalli. "Decentralized detection in sensor networks," *IEEE Trans. Signal Processing* **51** (2) : 407–416 (2003).

[11] F. R. K. Chung. *Spectral Graph Theory*, Regional Conference Series in Mathematics **92**. Providence, RI: American Mathematical Society (1994).

[12] J. A. Costa and A. O. Hero III. "Geodesic entropic graphs for dimension and entropy estimation in manifold learning," *IEEE Trans. Signal Processing* **52** (8) : 2210–2221 (2004).

[13] R. Diestel. *Graph Theory*. New York: Springer-Verlag, Inc. (2000).

[14] Doctor Web. "Doctor Web exposes 550 000 strong Mac botnet", 4 April 2012, accessed 3 September 2012 ⟨http://news.drweb.com/show/?i=2341⟩.

[15] W. E. Donath and A. J. Hoffman. "Lower bounds for the partitioning of graphs," *IBM J. Res. Development* **17** : 420–425 (1973).

[16] P. ERDŐS and A. RÉNYI, "On the evolution of random graphs," *Pubs. Mathematical Institute of the Hungarian Academy of Sciences* **5** : 17–61 (1960).

[17] J. P. FERRY, D. LO, S .T. AHEARN, and A. M. PHILLIPS. "Network detection theory," in *Mathematical Methods in Counterterrorism*, eds. N. MEMON et al., pp. 161–181, Vienna: Springer (2009).

[18] M. FIEDLER. "Algebraic connectivity of graphs," *Czech. Math. J.* **23** (2) : 298–305 (1973).

[19] ———. "A property of eigenvectors of non-negative symmetric matrices and its application to graph theory," *Czech. Math. J.* **25** : 619–633 (1975).

[20] S. FORTUNATO and M. BARTHÉLEMY. "Resolution limit in community detection," *PNAS* **104** (1) : 36–41 (2007).

[21] S. FORTUNATO. "Community detection in graphs," *Physics Reports* **486** : 75–174 (2010).

[22] C. GODSIL and G. ROYLE. *Algebraic Graph Theory*. New York: Springer-Verlag, Inc. (2001).

[23] GOOGLE. "The technology behind Google's great results," Accessed 3 September 2012 ⟨http://www.google.com/onceuponatime/technology/pigeonrank.html⟩.

[24] K. G. GURUHARSHA et al. "A protein complex network of *Drosophila melanogaster*," *Cell* **147** (3) : 690–703 (2011). Accessed 3 September 2012 ⟨http://www.sciencedirect.com/science/article/pii/S0092867411010804⟩.

[25] M. O. JACKSON. *Social and Economic Networks*, Princeton U. Press (2008).

[26] J. JONAS and J. HARPER. "Effective counterterrorism and the limited role of predictive data mining," *Policy Analysis* **584**. Cato Institute (2006).

[27] S. KAR, S. ALDOSARI, and J. F. MOURA. "Topology for distributed inference on graphs," *IEEE Trans. Signal Processing* **56** (6) : 2609–2613 (2008).

[28] J. M. KEYNES. *A Treatise on Probability*. London: Macmillan and Co. (1921).

[29] D. KOLLER and N. FRIEDMAN. *Probabilistic Graphical Models*. Cambridge, MA: MIT Press (2009).

[30] V. E. KREBS. "Uncloaking terrorist networks," *First Monday* **7** (4) (2002).

[31] J. M. KUMPULA, J. SARAMÄKI, K. KASKI, and J. KERTÉSZ. "Limited resolution in complex network community detection with Potts model approach," *Eur. Phys. J. B* **56** : 41–45 (2007).

[32] J. LESKOVEC, K. J. LANG, and M. MAHONEY. "Empirical comparison of algorithms for network community detection," in *Proc. 19th Intl. Conf. World Wide Web (WWW'10)*. Raleigh, NC, pp. 631–640 (2010).

[33] B. A. MILLER, M. S. BEARD, and N. T. BLISS. "Eigenspace Analysis for Threat Detection in Social Networks," in *Proc. 14th Intl. Conf. Informat. Fusion (FUSION)*. Chicago, IL (2011).

[34] B. A. MILLER, N. T. BLISS, and P. J. WOLFE. "Toward signal processing theory for graphs and other non-Euclidean data," in *Proc. IEEE Intl. Conf. Acoustics, Speech and Signal Processing*, pp. 5414–5417 (2010).

[35] ———. "Subgraph detection using eigenvector L_1 norms," in *Proc. 2010 Neural Information Processing Systems (NIPS)*. Vancouver, Canada (2010).

[36] B. MOHAR. "The Laplacian Spectrum of Graphs," in *Graph Theory, Combinatorics, and Applications*, **2**, eds. Y. ALAVI, G. CHARTRAND, O. R. OELLERMANN, and A. J. SCHWENK. New York: Wiley, pp. 871–898 (1991).

[37] R. R. NADAKUDITI and M. E. J. NEWMAN. "Graph spectra and the detectability of community structure in networks," *Phys. Rev. Lett.* **108**, 188701 (2012).

[38] M. E. J. NEWMAN. "Finding community structure in networks using the eigenvectors of matrices," *Phys. Rev. E*, **74** (3) (2006).

[39] ———. *Networks: An Introduction*.p Oxford U. Press (2010).

[40] S. PHILIPS, E. K. KAO, M. YEE, and C. C. ANDERSON. "Detecting activity-based communities using dynamic membership propagation," in *Proc. IEEE Intl. Conf. Acoustics, Speech and Signal Processing (ICASSP)*. Kyoto, Japan (2012).

[41] J. W. PERRY et al. *Protecting Individual Privacy in the Struggle Against Terrorists: A Framework for Program Assessment*. The National Academies Press (2008).

[42] A. POTHEN, H. SIMON, and K.-P. LIOU. "Partitioning sparse matrices with eigenvectors of graphs," *SIAM J. Matrix Anal. Appl.* **11** : 430–45 (1990).

[43] M. SAGEMAN. *Understanding Terror Networks*. Philadelphia, PA: U. Pennsylvania Press (2004).

[44] A. SHAMIR. "A survey on mesh segmentation techniques," *Computer Graphics Forum* **27** (6) : 1539–1556 (2008). Accessed 3 September 2012 ⟨http://www.faculty.idc.ac.il/arik/site/mesh-segment.asp⟩.

[45] J. SHI and J. MALIK. "Normalized cuts and image segmentation," *IEEE Trans. Pattern Anal. Mach. Intell.* **22** (8) : 888-905 (2000).

[46] S. T. SMITH, A. SILBERFARB, S. PHILIPS, E. K. KAO, and C. C. ANDERSON. "Network Discovery Using Wide-Area Surveillance Data," in *Proc. 14th Intl. Conf. Informat. Fusion (FUSION)*. Chicago, IL (2011).

[47] S. T. SMITH, S. PHILIPS, and E. K. KAO. "Harmonic space-time threat propagation for graph detection," in *Proc. IEEE Intl. Conf. Acoustics, Speech and Signal Processing (ICASSP)*. Kyoto, Japan (2012).

[48] TELEGEOGRAPHY. "Global Traffic Map 2010," PriMetrica, Inc. Accessed 3 September 2012 ⟨http://www.telegeography.com/telecom-maps/global-traffic-map/index.html⟩.

[49] R. TRINQUIER. *Modern Warfare: A French View of Counterinsurgency*. Westport, CT: Praeger Security International (2006).

[50] UNITED STATES ARMY. *Counterinsurgency: Field Manual 3-24*, Appendix B. Washington: Government Printing Office (2006).

[51] H. L. VAN TREES. *Detection, Estimation, and Modulation Theory*, Part 1. New York: John Wiley and Sons, Inc. (1968).

[52] U. VON LUXBURG, O. BOUSQUET, and M. BELKIN. "Limits of spectral clustering," in *Advances in Neural Information Processing Systems* **17**, eds. L. K. SAUL, Y. WEISS, and L. BOTTOU. Cambridge, MA: MIT Press (2005).

[53] D. J. WATTS. "Networks, dynamics, and the small-world phenomenon," *American Journal of Sociology* **13** (2) : 493–527 (1999).

[54] Y. WEISS. "Segmentation using eigenvectors: A unifying view," in *Proc. of the Intl. Conf. Computer Vision* **2** : 975 (1999).

[55] S. WHITE and P. SMYTH. "A spectral clustering approach to finding communities in graphs," in *Proc. 5th SIAM Intl. Conf. Data Mining*, eds. H. KARGUPTA, J. SRIVASTAVA, C. KAMATH, and A. GOODMAN. Philadelphia PA, pp. 76–84 (2005.)

[56] H. WOLKOWICZ and Q. ZHAO. "Semidefinite programming relaxations for the graph partitioning problem," *Discrete Applied Mathematics* **96–97** : 461–479 (1999).

[57] J. XU and H. CHEN. "The topology of dark networks," *Comm. ACM* **51** (10) : 58–65 (2008).

[58] L. YANG. "Data Embedding Research," Western Michigan University. Accessed 3 September 2012 ⟨http://www.cs.wmich.edu/~yang/research/dembed/⟩.

[59] M. J. YEE, S. PHILIPS, G. R. CONDON, P. B. JONES, E. K. KAO, S. T. SMITH, C. C. ANDERSON, and F. R. WAUGH. "Network discovery with multi-source intelligence, surveillance, and reconnaissance," *Lincoln Laboratory J.*, to appear.